SCTE•ISBE CABLE-TEC EXPO® 2020
VIRTUAL EXPERIENCE » OCTOBER 12-15

2020 Fall Technical Forum
SCTE·ISBE • NCTA • CABLELABS®

Connectivity and COVID-19: Maintaining QoE During a Crisis

Technical Paper prepared for SCTE•ISBE by

William McFarland
CTO
Plume
290 California Ave., Palo Alto, CA 94306
650-823-6315
Bill@Plume.com

Table of Contents

List of Figures

1. Introduction

Few events in modern times have had as large an impact on society as COVID-19. Across the globe, people have been encouraged to work- and school-from-home. In many cases, this was mandated by governments by shelter-in-place orders –and for some, this continues to be the 'new norm'. Not surprisingly, this has resulted in a huge shift in network usage patterns. Several trends are apparent in the data that will be presented in this paper:

- Large increases in people working on the internet at home during the workdays
- Dramatic increases in the number of hours that devices are active on the home network, particularly computers, phones, and entertainment devices
- The workday is spreading into the evening hours as families try to juggle work and family commitments

Remarkably, wired internet access networks have held up to the added load well. **Figure 1** shows the download and upload speeds as recorded at more than 15 million households that are managed by the Plume Cloud. The delivered speeds, as measured by Plume, are remarkably consistent before and after COVID-19 caused the large increase in working-from-home. In fact, average speeds even trended *upwards* across the time due to subscribers upgrading to higher tier services to meet their expanded needs. At the onset of the pandemic many predicted that broadband service providers would struggle to meet higher and changing bandwidth demands–many OTT TV providers *even* reduced streaming quality to assist–however, Plume data shows that the networks coped admirably,

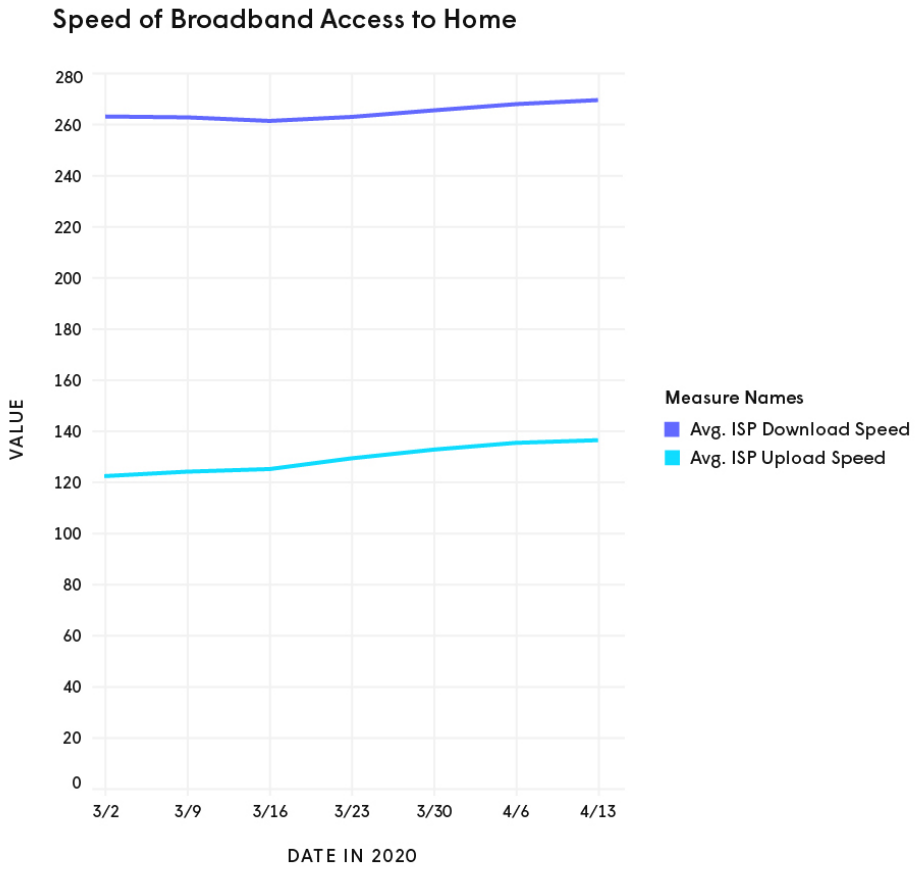

Figure 1 - Speed of Broadband Access to Homes Before and After COVID-19

While service providers have generally been able to keep up with the added load in delivering service to the edge of the customer's home, the story within the home could be much different. The majority of devices in homes today connect over Wi-Fi. Wi-Fi always has the fundamental problem of being a shared, and oftentimes best-effort medium, so handling increased load can be particularly challenging. **Figure 2** shows the dramatic increase in load from three key categories of devices: phones, computers, and entertainment devices (set top boxes, TVs, and gaming consoles). The increases are shown for both the amount of data consumed by these devices, as well as the number of minutes they were active on the network. In the stacked graph, the values are normalized to highlight the percent change from before to after COVID-19.

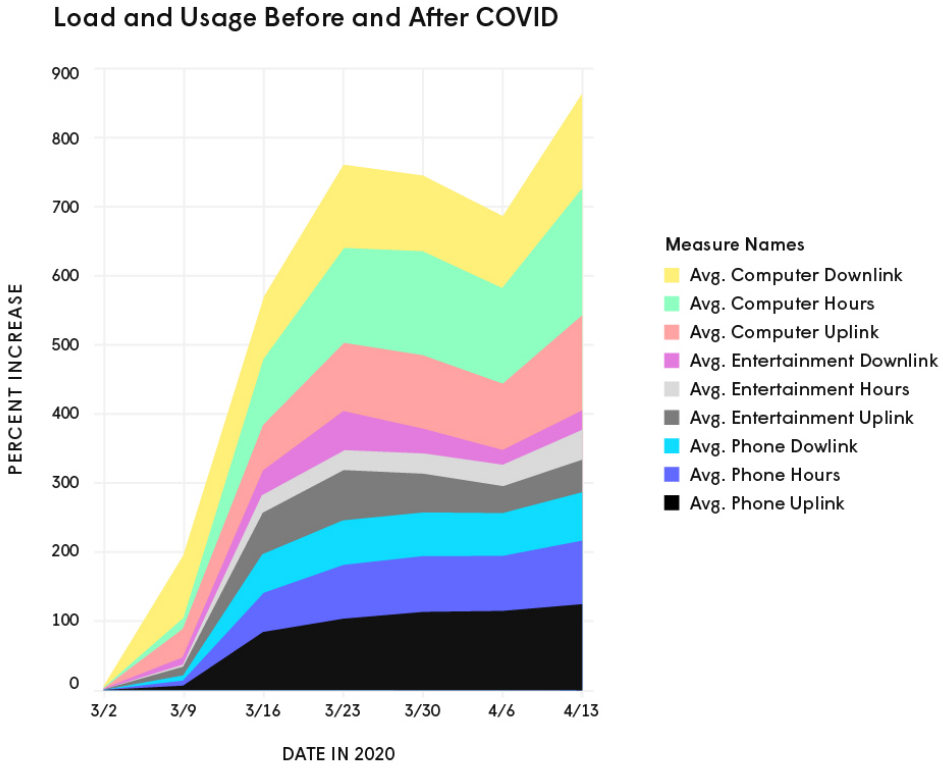

Figure 2 - Load and Usage Before and After COVID-19

The key takeaway from **figure 2** is that inside the home, a high degree of "wireless chaos" has been triggered by COVID-19. The remainder of this paper looks at this phenomenon in more detail, and describes what service providers can do to overcome this chaos and provide their customers with a highquality of *experience*.

2. The Effect of COVID-19 on User Behavior

As pointed out in the introduction, COVID-19 has changed network usage patterns significantly. This section examines these changes in more detail. While it is impossible to predict for sure, it is likely that

some of this behavioral change will persist long after COVID-19 has receded. If it recedes at all. For example, numerous companies, in the wake of the surge in working at home with COVID-19, have announced permanent acceptance of working from home. It is likely that many employees will take advantage of these offers. While this paper focuses on the changes before and after the onset of COVID-19, it is likely that the new behaviors, and the required adaptations by service providers for them, will be long term.

2.1. Changes in Working at Home

The most immediate effect of COVID-19 is an increase in the number of people working from home. Plume was able to analyze the data collected by Wi-Fi networks that it manages for service providers around the world to identify the change in the number of people who are working-from-home. The data analyzed included the type of devices, amount of data, amount of active time on the network, and domains accessed. Combining these observations, Plume is able to gauge homes in which at least one person is working from home during the 9am to 5pm weekday period. **Figure 3** shows the result for fourteen major metropolitan areas in the US:

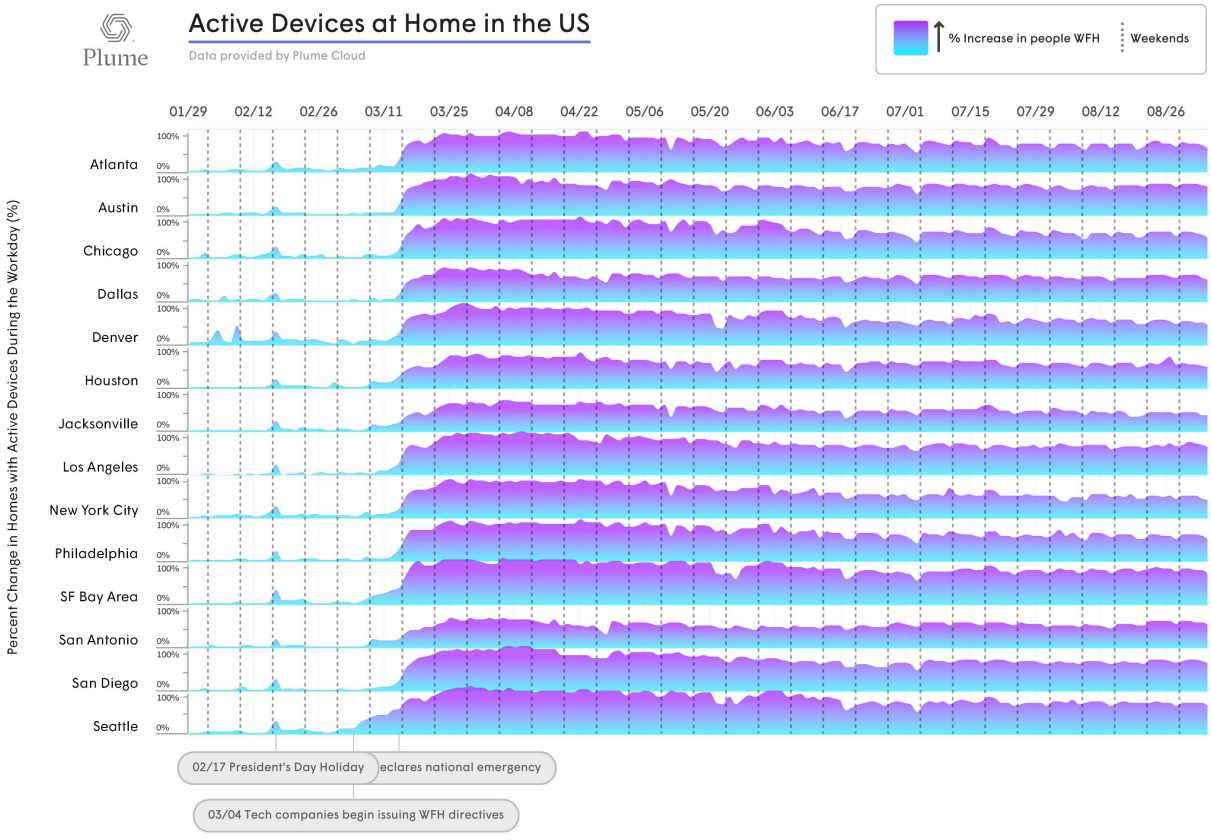

Figure 3 - Increase in Working From Home in the US

While the reaction across cities is relatively consistent, some details are instructive. The rise in working from home (WFH) occurred first in Seattle, the first city to see a significant number of cases, and the first

city to declare a shutdown. The San Francisco (SF) Bay Area had among the highest levels of extended WFH, perhaps enabled by its high percentage of tech workers. Cities in certain regions of the country (e.g. San Antonio and Jacksonville) have had somewhat lower levels of WFH. And we see some aborted attempts at loosening, followed by returns to higher levels of WFH.

Plume manages a large number of Wi-Fi networks in Canada and Europe as well as the USA. While the trends are similar, we can see a stronger, steadier downward trend in working-from-home in several Canadian cities and European countries **(Figures 4 and 5)**. These correlate with countries that have been better able to control their COVID-19 case counts.

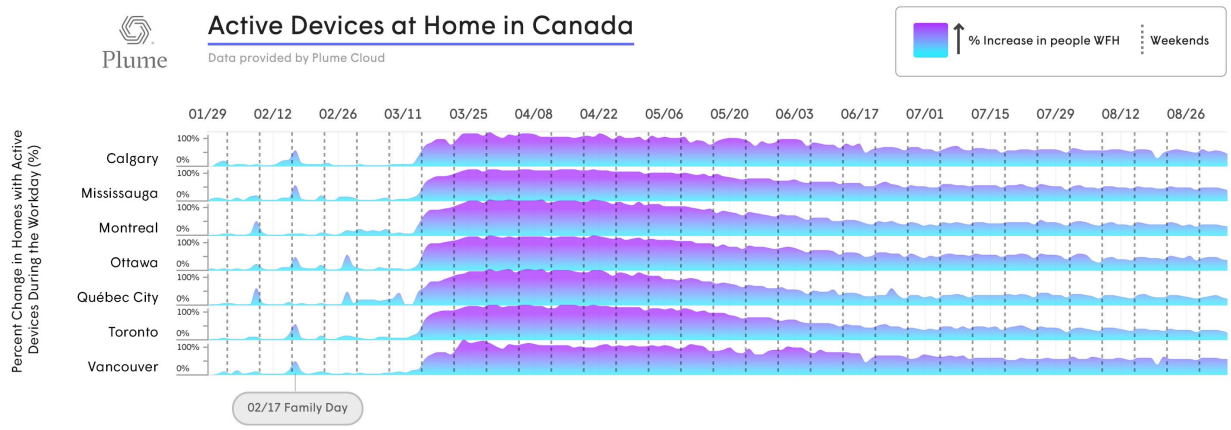

Figure 4 - Working From Home in Canada

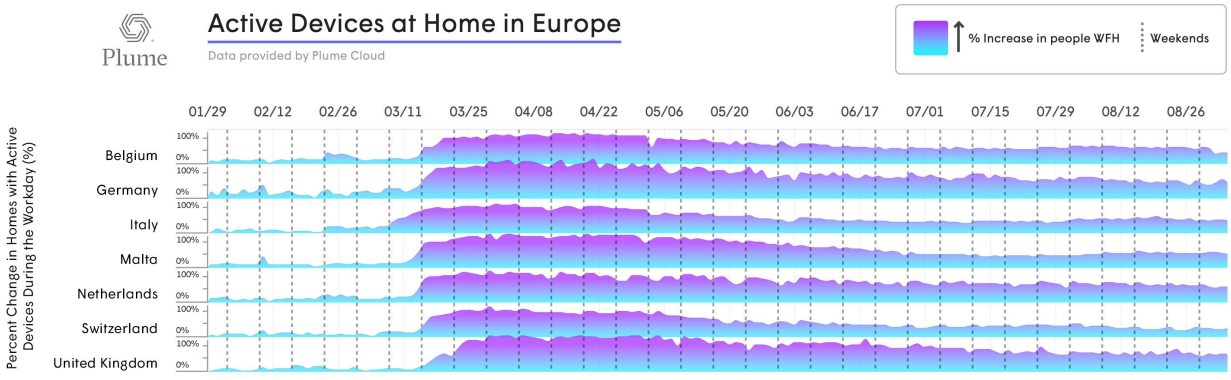

Figure 5 - Working From Home in the EU

2.2. Changes in Usage

Increased WFH is the first link in the chain between COVID-19 and increased in-home network distress. The change in usage of devices is the second. Devices don't readily identify their device type when they connect to a Wi-Fi network. But Plume is able to use a variety of factors including DHCP requests, UPNP transactions, DNS requests, Host Names, and user agent (browser) transactions to identify device types. This is based on machine learning across tens of millions of homes that are operated by Plume. After identifying the device type, Plume is able to observe the amount of data consume by the different device types, as well as the number of minutes that different device types are active on the network. **Figure 6** shows Plume's observations regarding the number of active minutes by device type across homes in the US. The "Computers" category includes both desktop and laptop computers, and the "Entertainment" category includes set top boxes, TVs, and gaming consoles.

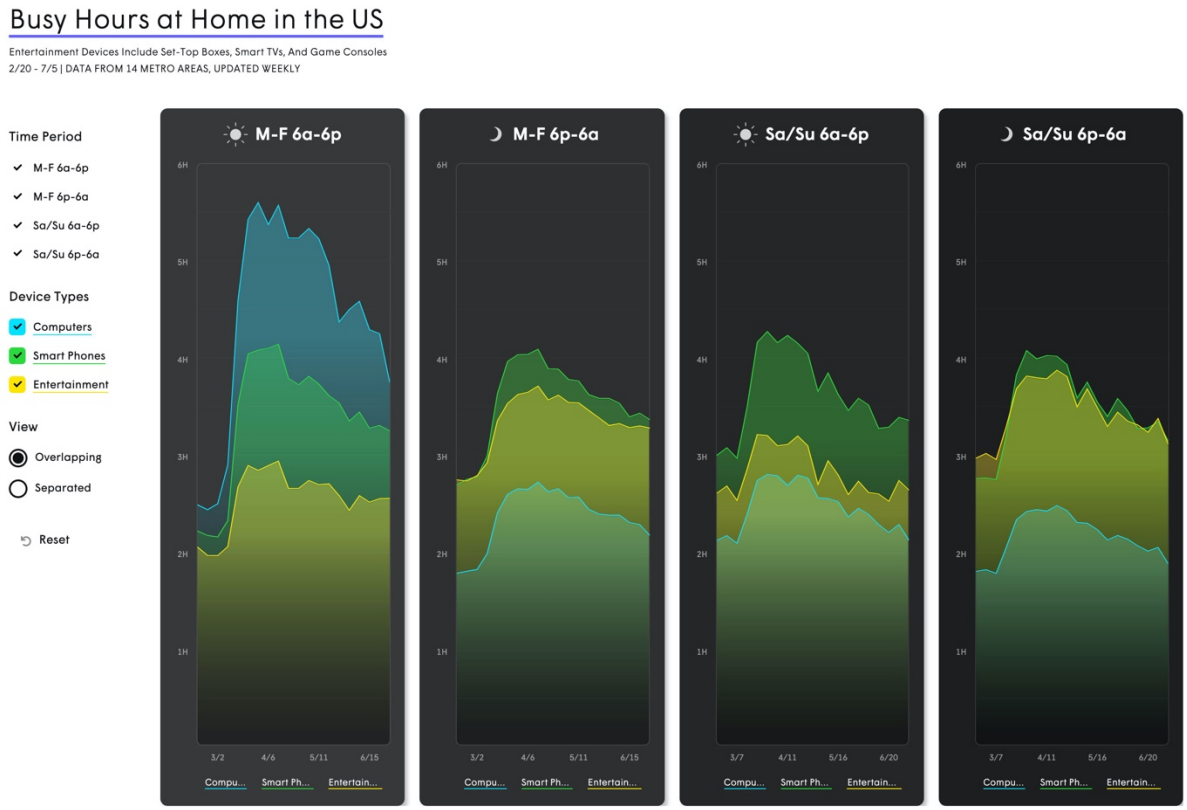

Figure 6 - Hours of Active Time by Device Type

By breaking the usage into days of the week and times of the day, several interesting before/after COVID-19 conclusions can be drawn. Unsurprisingly, the most dramatic increase in usage was for computers during the weekday work times. However, significant increases in phone and entertainment usage are seen in the evenings and on the weekends. This is likely due to the reduced entertainment options with social distancing. However, computer usage has grown significantly in the evenings and on weekends, indicating some spreading of working hours, as well as increased use of computers for entertainment.

3. The Effect of COVID-19 on Networks and Technologies to Compensate

The third link in the chain from cause to effect is how changes in usage due to COVID-19 impact the networks in people's homes. This section focuses on five of the most important effects. While these effects make the provision of adequate connectivity to provide high quality of experience more difficult, there are technologies that can compensate for the added challenges. These technologies are presented side by side with the description of the problem.

3.1. Demands on Coverage, and Multi-AP Solutions

The most basic requirement for any Wi-Fi user is to be able to physically connect to the network. When people have trouble connecting in their own home, the problem is usually caused by inadequate coverage. Coverage is defined the extent to which a given home has adequate signal strength to and from all devices at every location in the home. Homes often have "dead spots," regions that are not covered well by Wi-Fi. Behavioral changes due to COVID-19 have exacerbated this problem. As multiple people in the home search out areas where they can work in privacy, they are more often trying to use regions of the house with poor coverage.

Since client devices typically transmit at lower power levels than the Access Points (APs), it is often the connection that comes back from the client device that is the limiting factor. Plume's networks observe the signal strength of transmissions coming from client devices, moving this information to the cloud. Based on analysis of this information in the cloud, it is possible to determine which homes have a coverage problem. For this analysis, Plume defines a home to be in coverage alarm if more than 25% of the client devices have a coverage issue, a coverage issue defined as a device with under -70dBm signal strength 50% of the time. -70dBm is the level at which data rates start to drop quickly, and clients start to spend a lot of time searching for better APs or roaming, making media flows less reliable. **Figure 7** below shows the percentage of homes that are in coverage alarm for a large North American service provider. As can be seen, more than 40% of homes that have only one Wi-Fi AP are in coverage alarm. However, by adding additional APs ("Pods"), that can be reduced close to 10%.

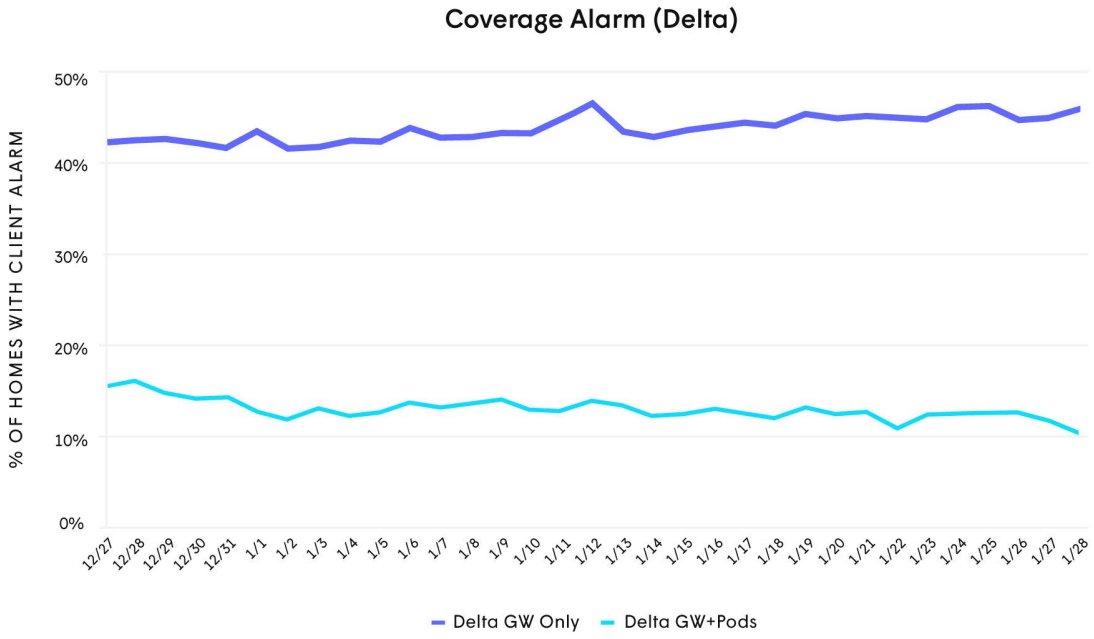

Figure 7 - Coverage Alarm vs. GW only or GW + additional APs

Typically, additional APs added into a home are connected wirelessly to each other. Few homes have Ethernet wires in locations that are appropriate for placing additional APs. Traditional Wi-Fi repeaters or mesh systems can solve the coverage problems just discussed, but they do so at the expense of throughput. These systems will utilize the same frequency channel for the backhaul connection between APs as is used to connect client devices to the AP (fronthaul). In fact, mesh systems often deploy the entire mesh, including all fronthaul and backhaul connections, on the same frequency channel. Such an arrangement suffers from self-interference in which transmissions on one hop in the network interfere with transmissions on other hops. For a two hop path, the throughput is more than halved, for a three or four hop path, the throughput is divided by more than 3 or 4 respectively.

This can be greatly improved if a different frequency channel can be used on each hop in the network. **Figure 8** shows a comparison between three approaches for serving a 20 meter connection in a typical home: direct connection to a powerful single AP, multiple hops through traditional repeaters or mesh system, and multiple hops through an optimized Adaptive Wi-Fi network. The difference in throughput achieved to the client is dramatic.

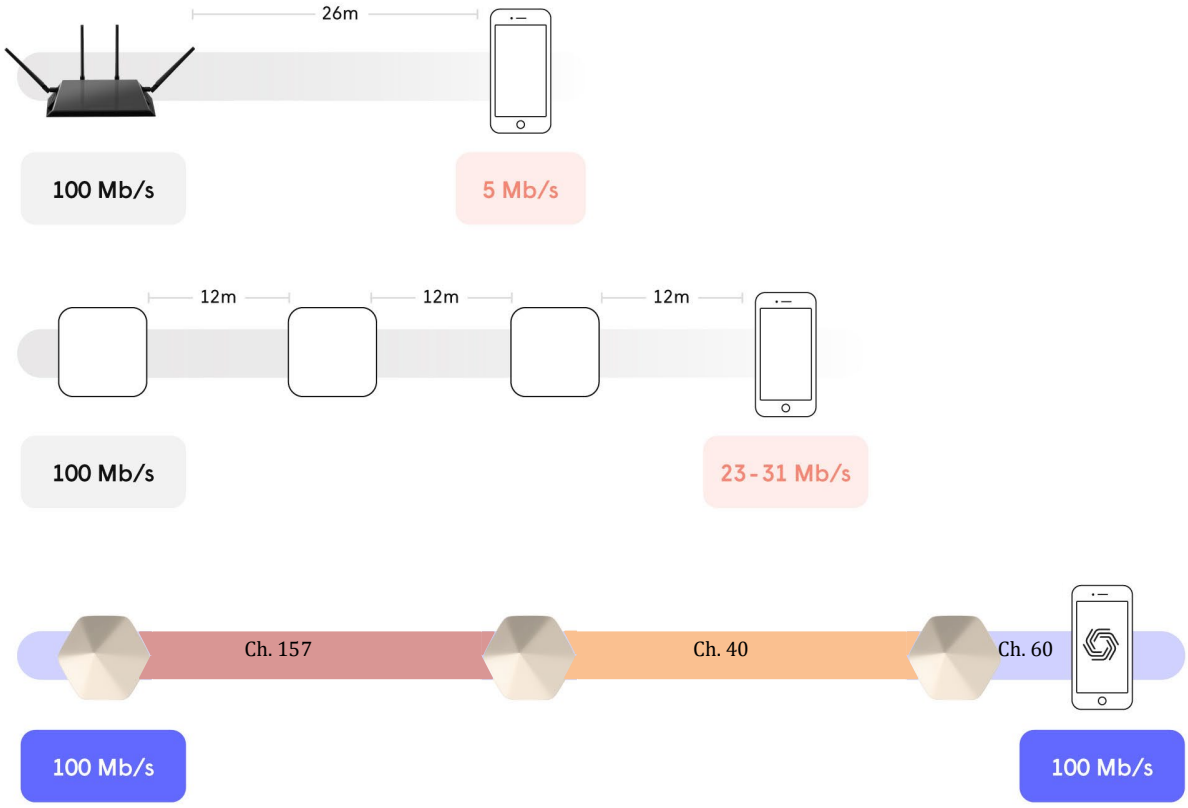

Figure 8 - Single AP vs Traditional Mesh vs Optimized Adaptive Wi-Fi

The complexity of choosing and configuring the optimized Adaptive Wi-Fi topology is why it is not often employed. However, modern cloud technology can enable sophisticated approaches. Rigorous optimization approaches, such as Mixed Integer Linear Programming (MILP), can be employed to maximize an objective function that includes throughput to individual clients, overall home system capacity, and fairness among the devices in the home. **Figure 9** shows a conceptual diagram of such an optimization system with its inputs and outputs. This approach is advantageous because the resulting topology deployed in a particular home will be optimal given the constraints of that particular home.

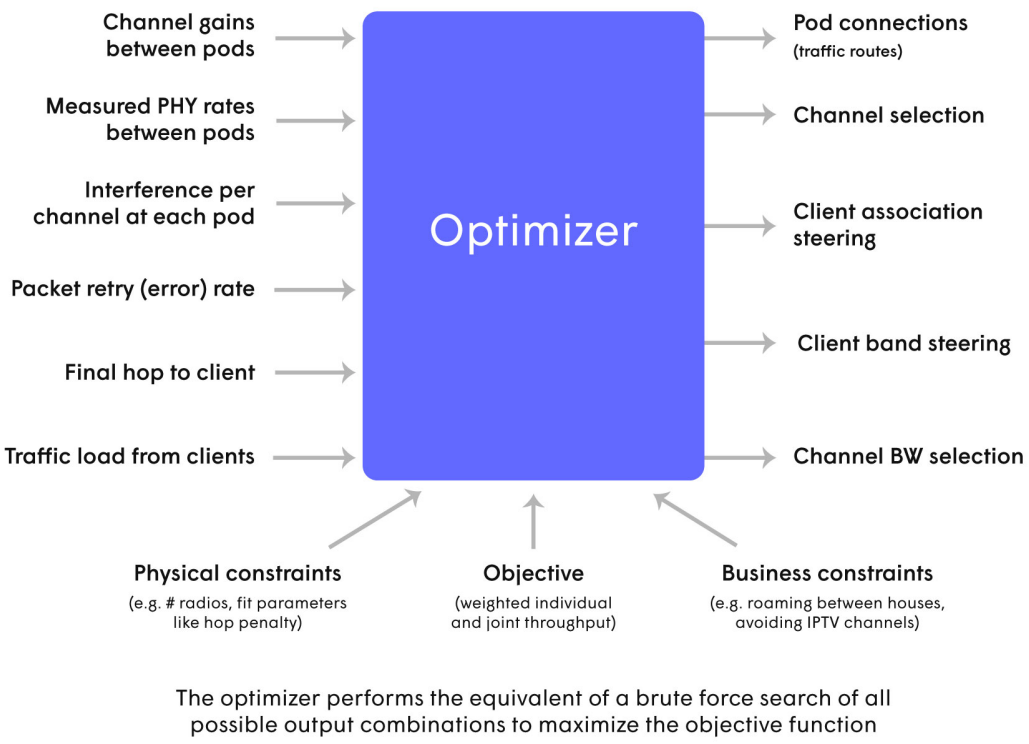

Figure 9 - Optimization System for Wi-Fi Networks

3.2. Increased Load, and Throughput Optimized Steering

The optimized Adaptive WIFI system presented in section 3.1 solves the problem of coverage while preserving optimal throughput. However, multi-AP systems in general bring the problem of where client devices connect into the network. Two problems are generally seen. First clients may be "sticky", remaining attached to one AP even as they move to the far side of the home. This has led some manufacturers to implement "sticky client steering" in which clients that are detected to be "stuck" on far away APs are kicked off that AP, forcing them to search for a closer AP for connection to the network. While this is better than nothing, the connections achieved in this way are not optimum. And second, even clients that dutifully shift to the closest AP are not necessarily connected in the optimum way. This is fundamental, as the client cannot know the complete pathway through the Wi-Fi network to the Internet. All the client can know is the signal strength from it to the APs it might connect to. To achieve the optimum connection requires understanding the speed of the complete connection, perhaps through multiple hops, from the client to the Internet for each option. [In response to SCTE editor comment.]

Figure 10 shows the desired outcome in client steering and AP topology. First, the optimizer eliminates excessive hops that would degrade performance. It does this for example by connecting APs directly to the GW AP as appropriate, when an extra hop to get through the network is not justified. Second, throughput optimized steering is used to move the client device to the AP where it will achieve the best

performance. The highest performance connection point for a client is often *not* the closest AP. As can be seen in the figure, connecting to the closest AP sometimes adds unnecessary additional hops, and can even increase the individual hop lengths in the connection.

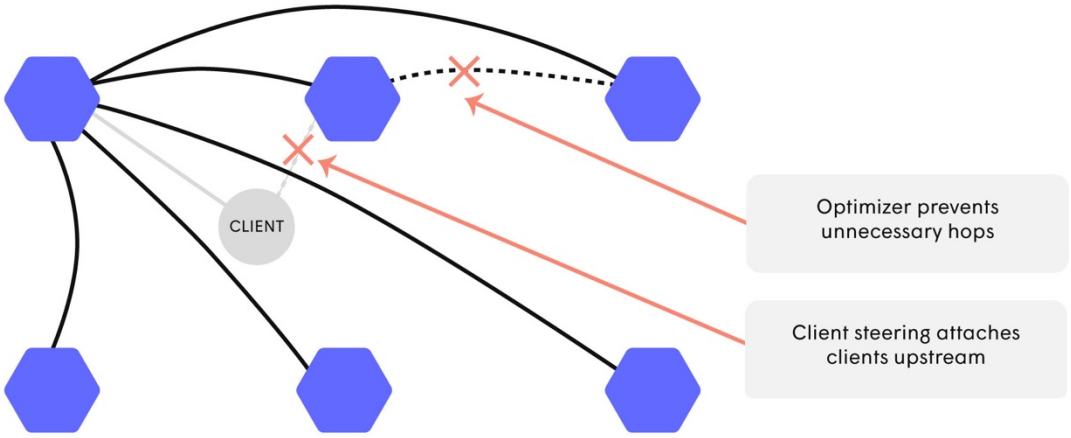

Figure 10 - Throughput Optimized Steering

The best location for client devices can be calculated as part of the optimization procedure. As clients move about the home and attach in inappropriate places, the results from the optimization can be used to identify where they would perform best, and client steering can be used to move them to the correct location.

3.3. Interference and MDU Joint Optimization

Increased usage due to COVID-19 also causes an increase in interference. Both interference (overlapping transmissions from neighbors) and congestion (self-interference from other devices within the home) are increased. While the increases occur across the board, these increases are particularly problematic in dense living environments such as Multi-Dwelling Units (MDUs), for example apartment complexes. **Figure 11** shows interference levels in suburban homes vs. urban MDUs as collected on Plume managed Wi-Fi networks. The graphs are histograms of the interference levels. The x-axis is the percentage of airtime consumed by the interference. For example, 33% interference indicates neighbors are consuming 33% of the available airtime, reducing throughput in the home in question by more than a third. The y-axis shows the percentage of homes/apartments in the particular interference bin of the histogram. 100% of homes surveyed are represented at some point in the histogram.

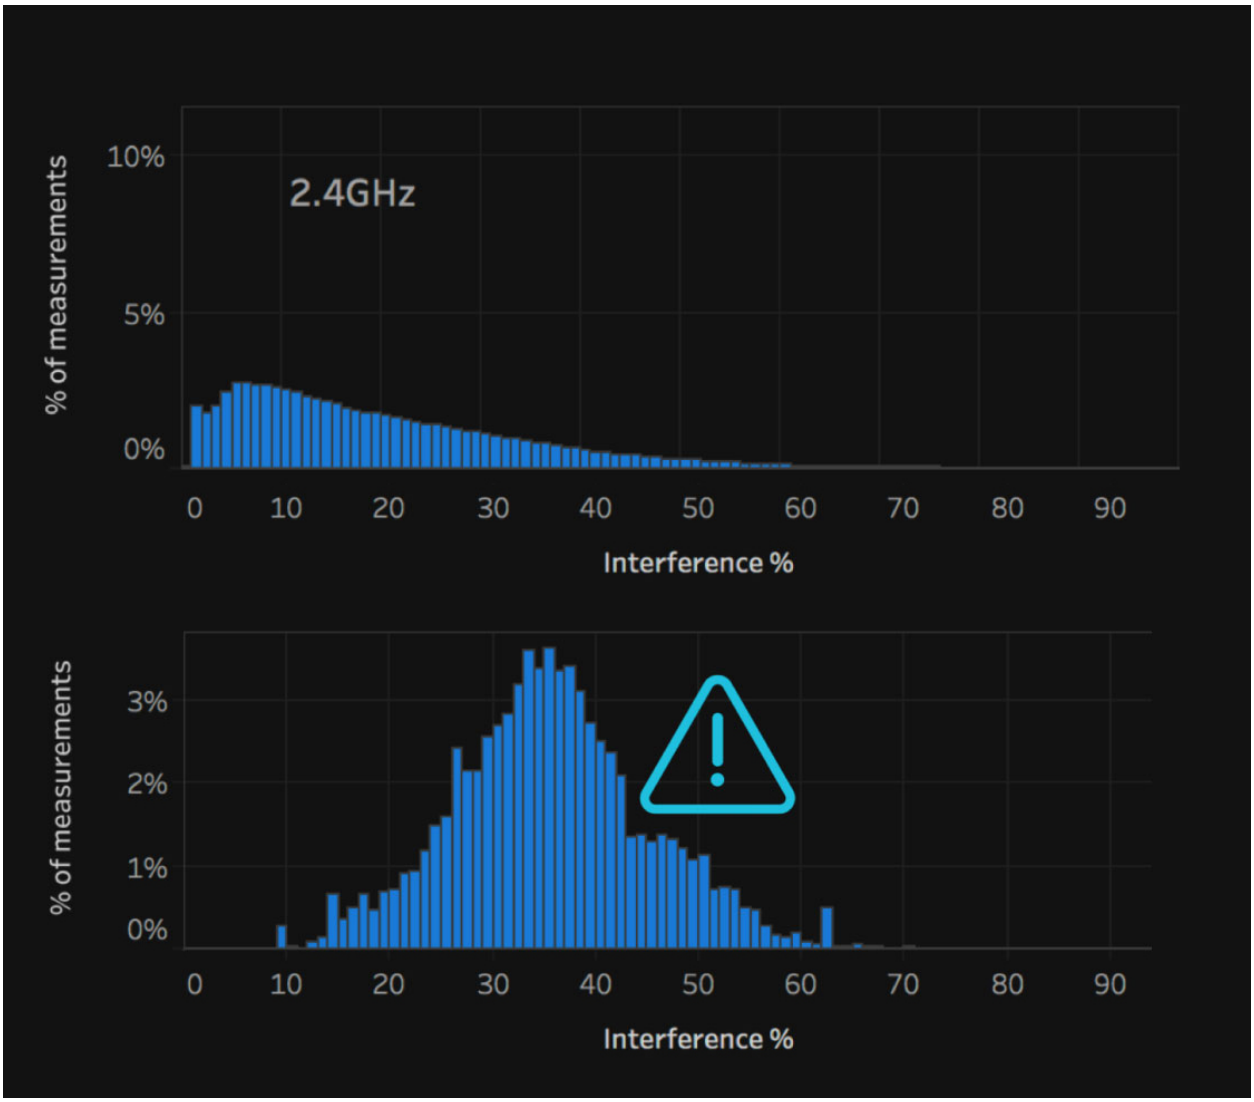

Figure 11 - Interference Histogram for Surburban Homes and Urban MDUs

As with previous COVID-19 related issues, the application of sophisticated cloud based processing can mitigate interference. In this case, the key is to consider the MDU as a whole. The algorithm begins with the APs reporting statistics about the neighboring APs they can see, along with interference levels and traffic loads. Based on the lists of neighboring APs, a clustering algorithm, run in the cloud, forms groups of APs that are tightly coupled to each other. The algorithm balances cluster size with the completeness of including all APs that interact with other APs in the cluster. Once the clusters have been formed, the optimizer can select frequency channels and bandwidths so as to maximize the performance of the entire cluster. **Figure 12** shows the result for a sample MDU.

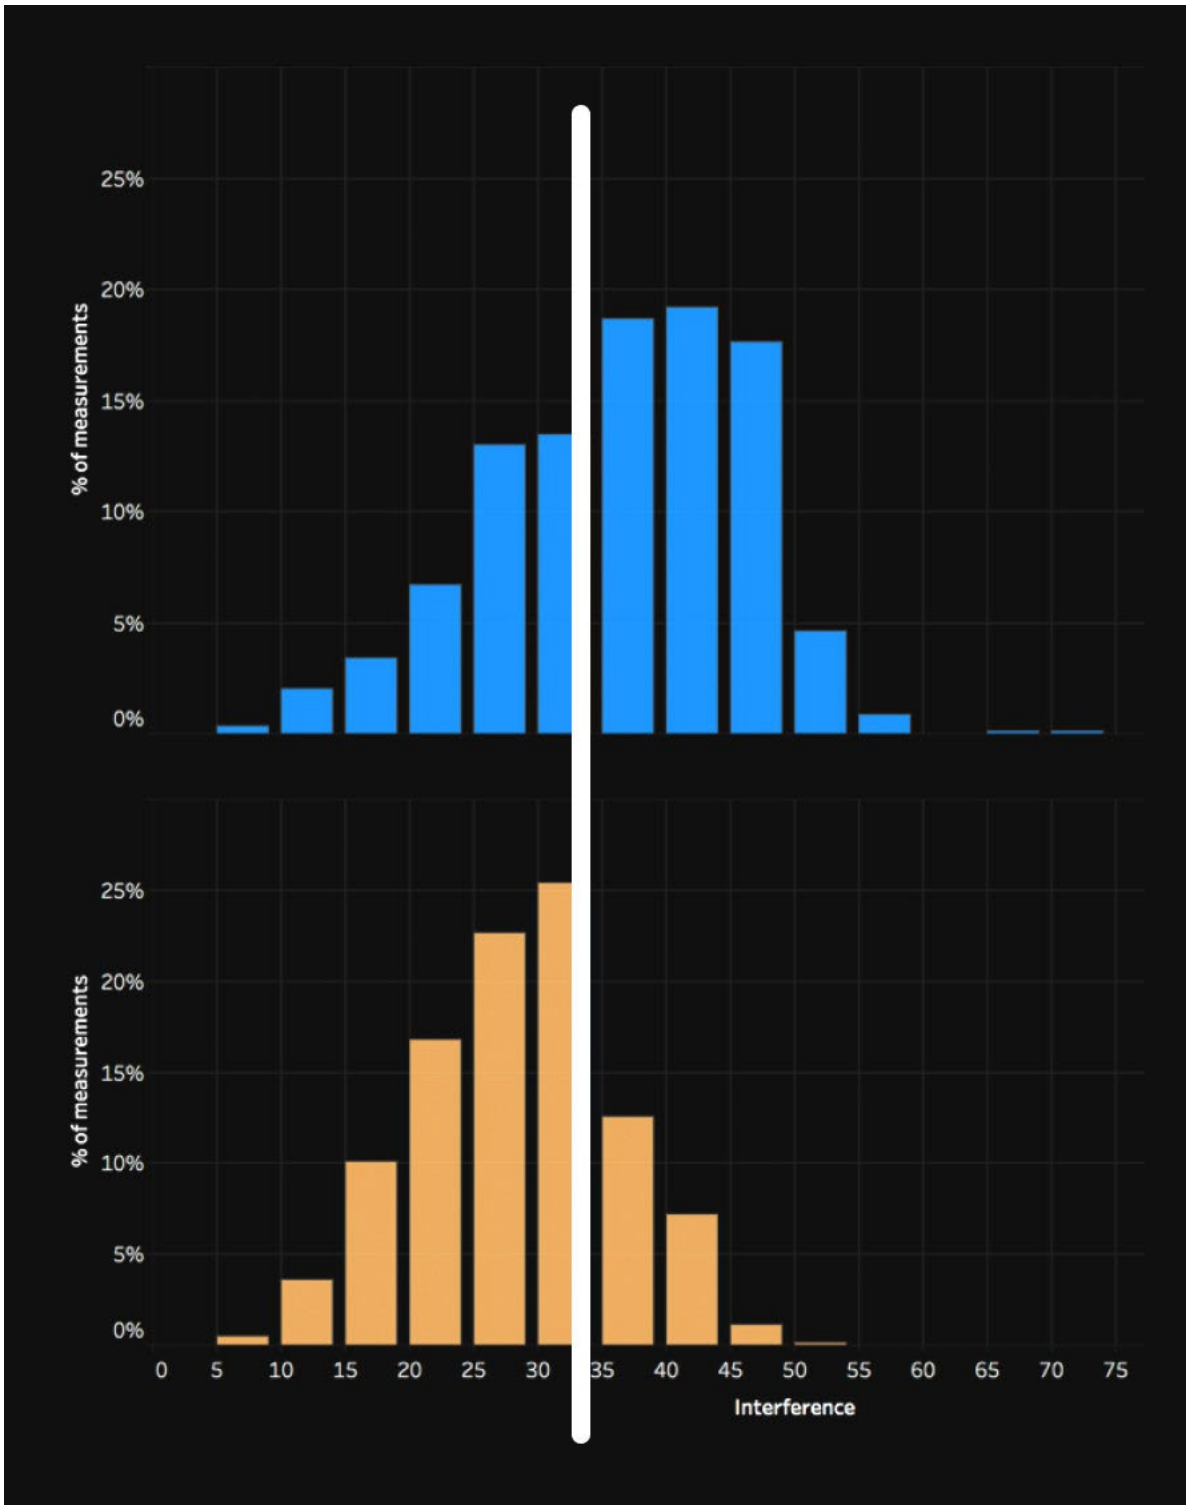

Figure 12 - Before and After MDU Joint Optimization Interference Histograms

The upper histogram shows the distribution of homes vs. interference levels before optimizing together as a cluster. The lower histogram shows the distribution after optimizing together as a cluster. The vertical line marks the point at which 33% of airtime in homes is consumed by interference. The joint optimization is able to significantly reduce the number of apartments experiencing interference above 33%, the point at which interference effects become significant.

3.4. Cyber Security Attacks and IoT Device Security

It has been hypothesized that criminals would take advantage of the fear, confusion, and increased internet use due to the COVID-19 pandemic. Sadly, it turns out this is true. **Figure 13** shows the increase in various types of threats before and after COVID-19, with several attack types doubling in frequency. In fact, across the period of this study 87% of the homes had some type of cyber security attack.

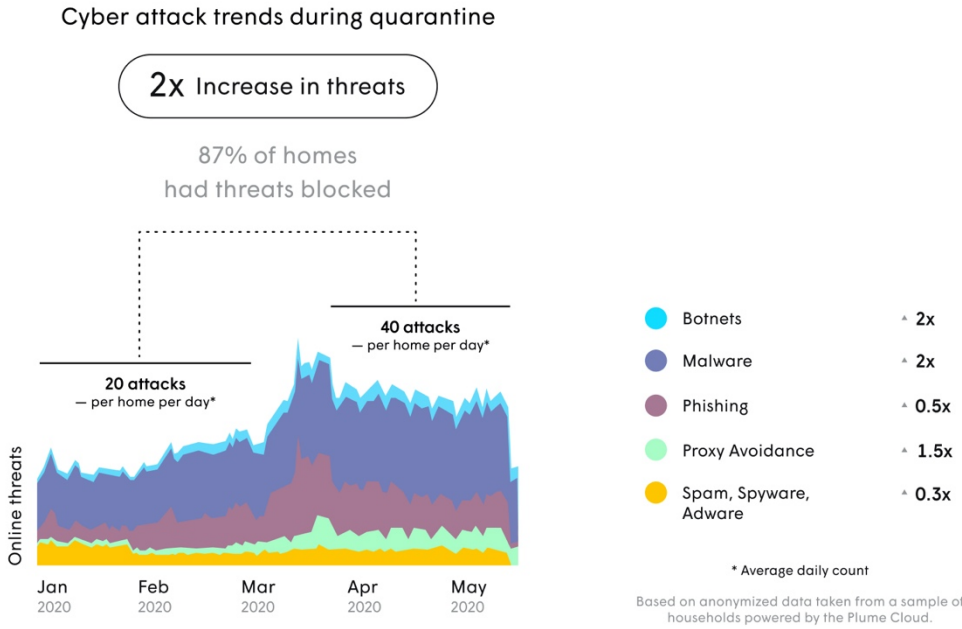

Figure 13 - Cyber Security Attacks Before and After COVID-19

To combat this, a robust cyber security system should be deployed. A system that includes a variety of layered protections will be most effective. Each of the following methods have their own advantages and weaknesses, but when taken together as a whole form an effective defense.

DNS query checking: Devices making an outbound connection to the internet often begin that transaction with a Domain Name Service (DNS) query to obtain the IP address for a given domain name. Domain names are generally logical and somewhat understandable, so it is relatively easy to maintain a list of high risk and low risk domains. A disadvantage is that creative cyber criminals avoid the use of domain names by creating viruses that use IP addresses immediately, without a domain name lookup.

IP address reputation checking: IP addresses can be checked similarly to domain names. The advantage is that this works well for incoming traffic as well as outgoing traffic, and it can cover cases in which viruses avoid DNS lookups. However, IP addresses change frequently, and maintaining a valid list of acceptable vs. unacceptable IP addresses is difficult.

Port checking and anomaly detection: Internet connections usually have an associated port number, which is frequently indicative of the application being used. Certain applications, such as ssh, should not be active on certain types of devices. Such basic rules can be augmented by doing machine learning (ML) based anomaly detection. The ML approach can be particularly powerful if data from millions of homes can be aggregated in the cloud, forming a robust dataset for training.

Tx/Rx Data pattern anomaly detection: The last line of defense is to observe the basic behavior of the device. For example, **Figure 14** shows a plot of the number of bytes transmitted and received by a Nest Camera in a given minute plotted as a two dimensional scatterplot. Regions that define different modes of operation can be identified, and behavior that does not fall into these regions can be flagged as anomalous, potentially caused by a virus or malware. This technique is particularly helpful for IoT devices, which are generally headless, and can't have aftermarket anti-virus software added. As with all ML based anomaly detection techniques, having a large training set is crucial. Cloud based management systems that aggregate data from millions of homes are beneficial for this approach.

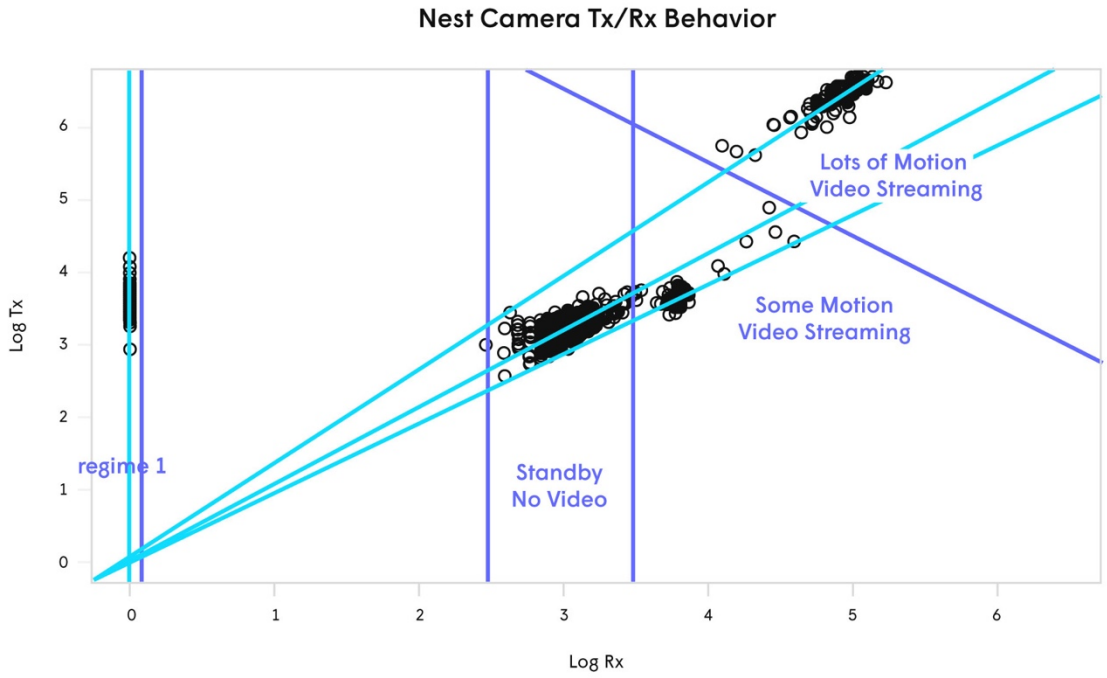

Figure 14 - Scatterplot of Tx and Rx Bytes per Minute for Nest Camera

Cyber Security at every node: Cyber security measures have traditionally been implemented only at the gateway or main router in a home. However, with the advent of extenders and repeaters, this invites the lateral movement of viruses or malware between devices in the home. Traffic that flows within the home is naturally routed over the shortest path, and this path may not include the gateway or router. And, many mobile devices leave the home and return, sometimes with a new virus on the device. Figure 15 shows the need to detect and stop viruses even for traffic flowing only through a repeater or extender.

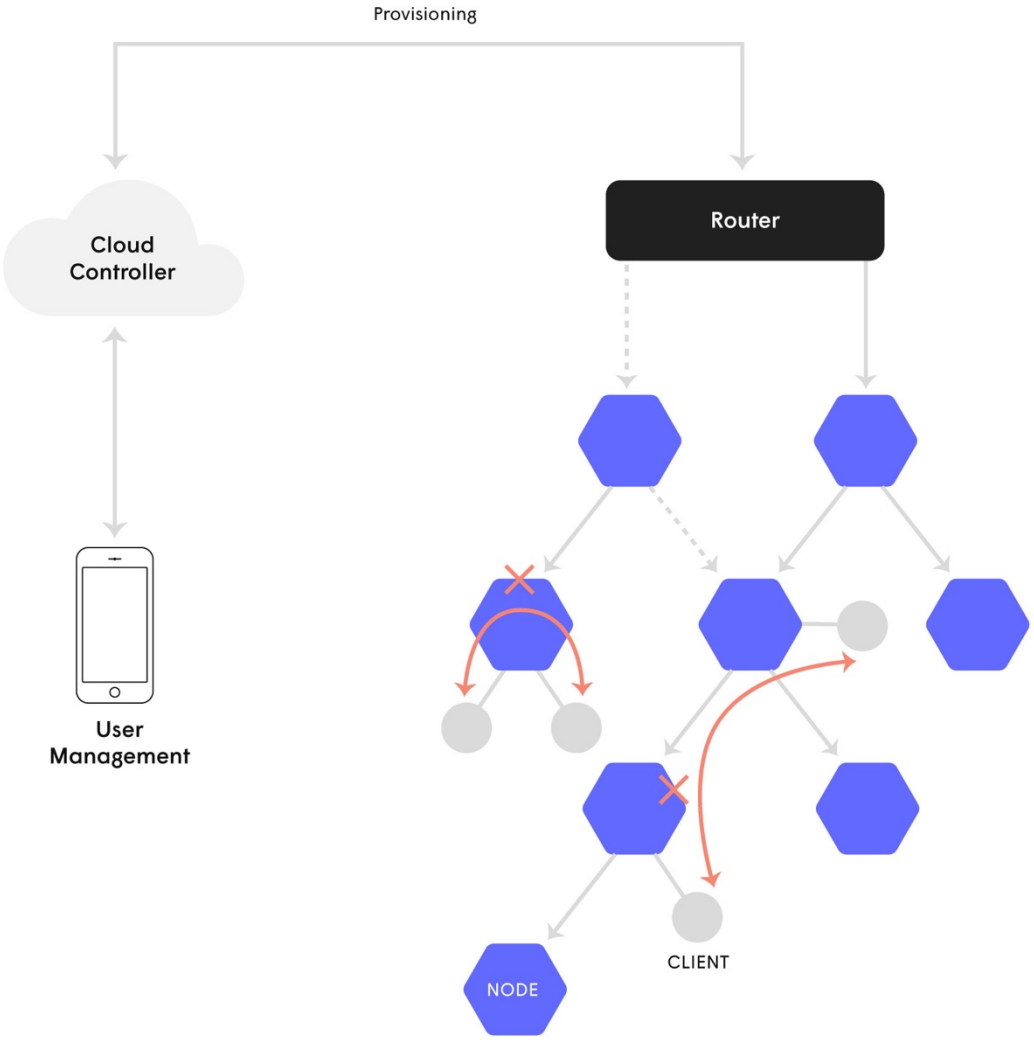

Figure 15 - Lateral Movement of Viruses and Security at Every Node

3.5. Traditional QoS vs. QoE

As COVID-19 puts more demands on networks, carriers are naturally concerned whether their networks are performing adequately to satisfy their customers. Quality of service (QoS) is the metric that has traditionally been used to indicate this. Typical Wi-Fi QoS metrics factor signal strength, data rate, and congestion/interference, to indicate the quality of the operation of the network. However, consumers are not actually interested in the performance of the network. Consumers are interested in the experience they have using services over those networks. To accurately reflect consumer's satisfaction, Quality of Experience (QoE) is a more effective metric than QoS. QoE starts with similar factors as QoS, but extends that to consider the devices and services that customers are using in a particular home. As an example, a customer with a Wi-Fi thermostat that receives reliable 1Mb/s service will be perfectly happy. On the other hand, a customer with a 4k set top box that receives 20 Mb/s service is likely to be unhappy. Along with factoring the type of device or service in question and their needs, the QoE metric factors far more conditions than just signal strength, data rate, and congestion. **Figure 16** shows the factors considered in a QoE algorithm, and how such a QoE score can be used.

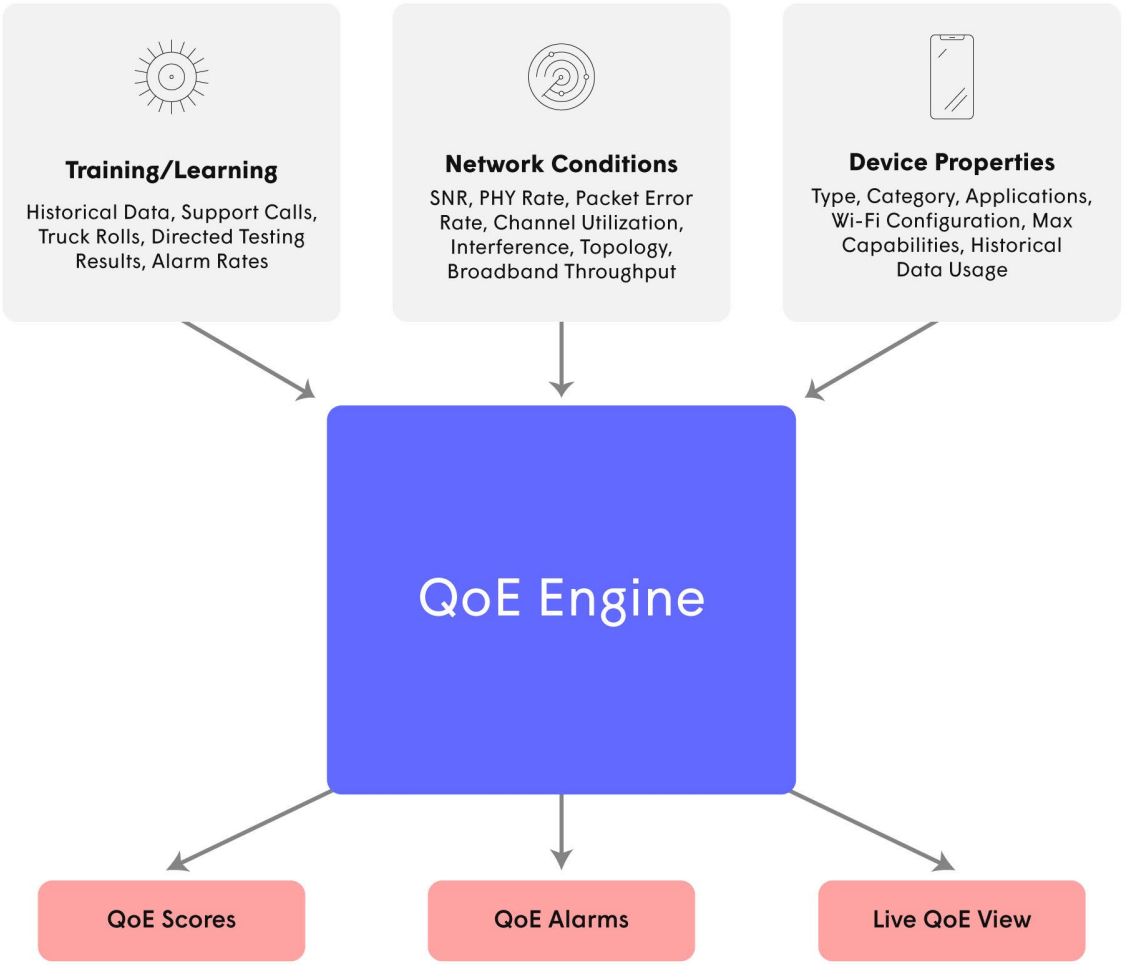

Figure 16 - QoE Factors and Outputs

The resulting difference between the QoS and QoE metrics can be compared by using traditional QoS metrics to identify devices that are performing poorly, comparing that set to devices identified using the more accurate QoE metric. **Figure 17** shows the statistics of such a comparison on networks operated by Plume. The Venn diagrams indicate that the different methods identify largely different sets of devices as requiring attention.

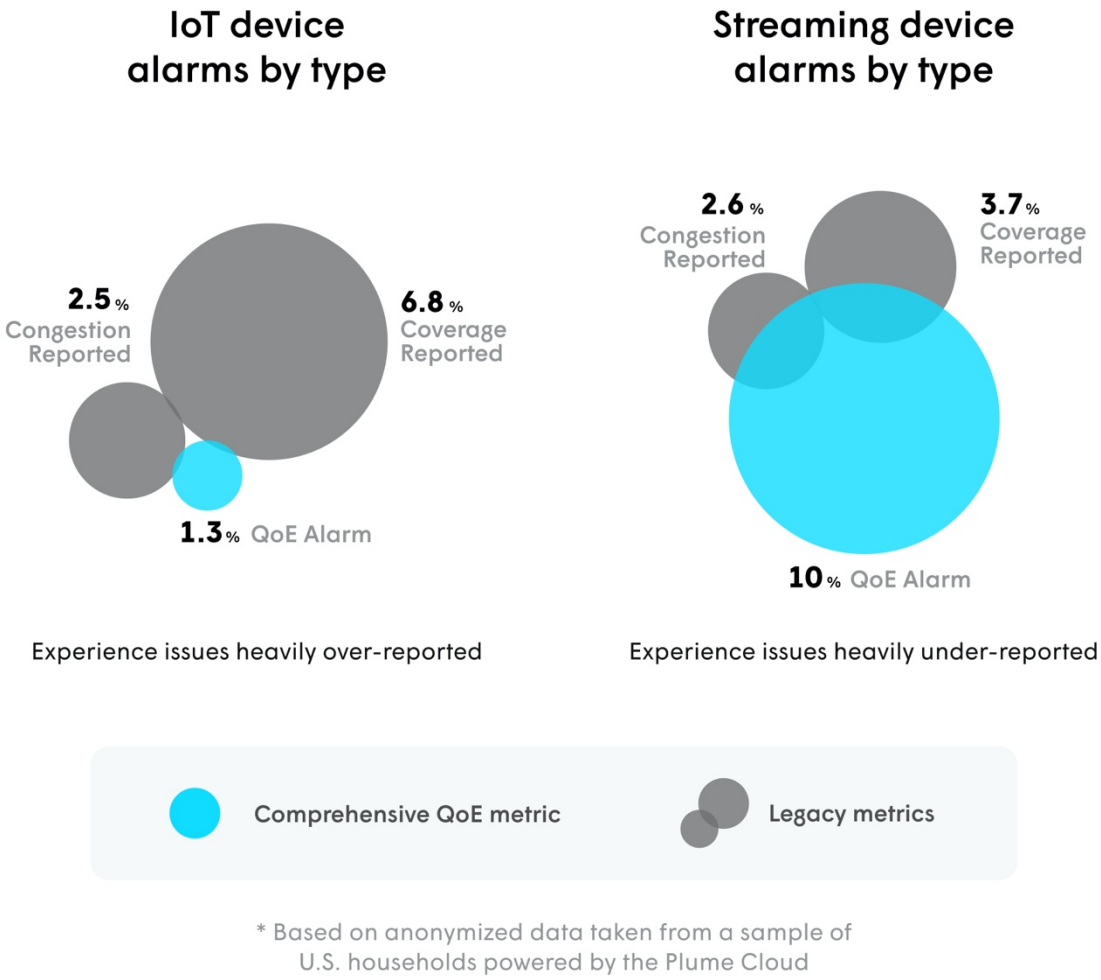

Figure 17 - Struggling Devices Identified by QoS and QoE

As could be predicted from the description above, too many IoT devices, which generally have lower networking requirements, tend to be identified as needing attention by traditional QoS approaches. This is significant, as service providers might spend too much money trying to fix those problems with additional repeaters or other steps. Conversely, streaming devices are under identified as needing attention by traditional QoS mechanisms. This also can be costly, as customers experience poor video quality and churn to another service provider.

4. The Effect of COVID-19 on Financials and How to Compensate

Luckily, COVID-19 is not having the devastating financial effect on service providers that it is having on other industries. However, the impact on carriers is still significant. New technologies can help service provider financials in at least two areas: the costs of customer support, and additional revenue from new services.

4.1. Proactive Support

As COVID-19 drives heavier use of home networks, and has consumers spending significant time on sensitive applications like teleconferencing, customer support becomes more heavily loaded. An innovative application of machine learning is to predict which customers are most likely to call in the next period of time. Such a capability can be used in a variety of ways:

- Direct proactive maintenance/care
- Send preventive email
- Understand what drives calls

ML cannot predict who is going to call perfectly, but it can do quite well. And, the algorithm can be adjusted to trade precision (the percentage of people identified by the algorithm who really would have called), and recall (the percentage captured on the list of all people who would have called). **Figure 18** shows results that Plume achieved on data from a North American deployment.

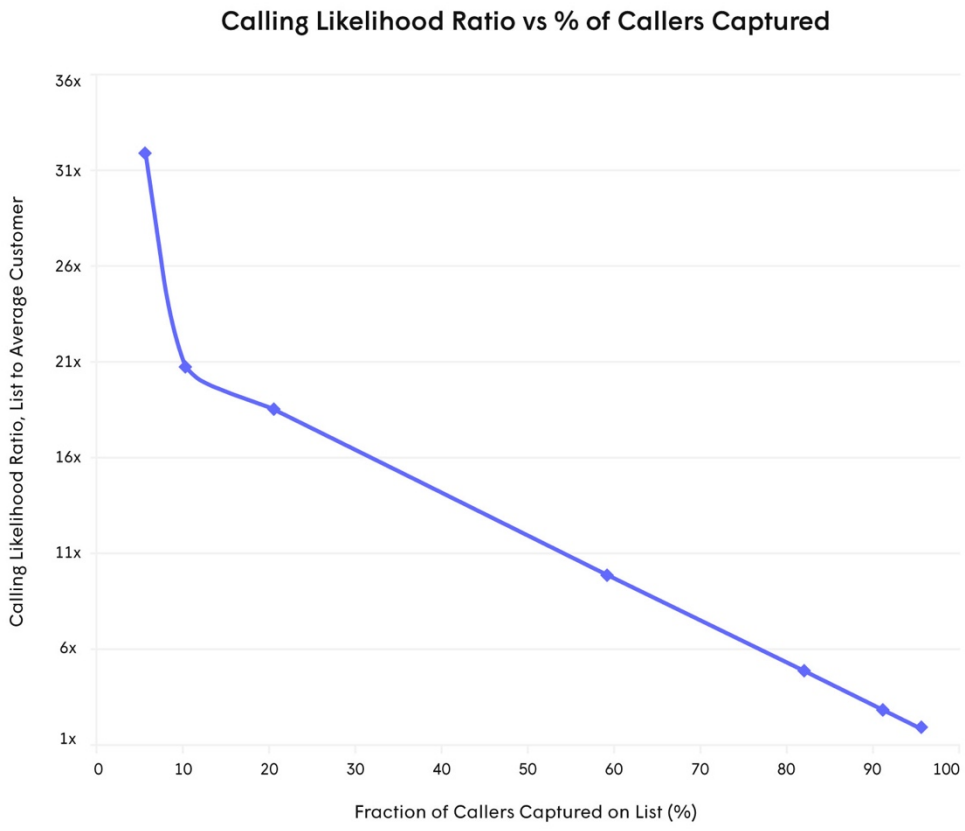

Figure 18 - Support Call Prediction Precision vs. Recall

Taking one particular point off the curve, it shows that we can form a list of people expected to call which includes 60% of people who actually will call (the x-axis), and the people on that list will be statistically 10x more likely to call than the population as a whole. The curve shows how these two values can be traded against each other.

4.2. Lost Revenue and Compensation with Additional Services

Another approach to make up for profit lost due to COVID-19 is to introduce additional services that users will pay a premium for. An innovative new service that can be offered using the in-home Wi-Fi network is motion detection. Motion in a home can be detected by watching the Wi-Fi signals between devices in a home. These signals will change in time if people are moving in the home due to changes in attenuation and reflection patterns. **Figure 19** illustrates the concept:

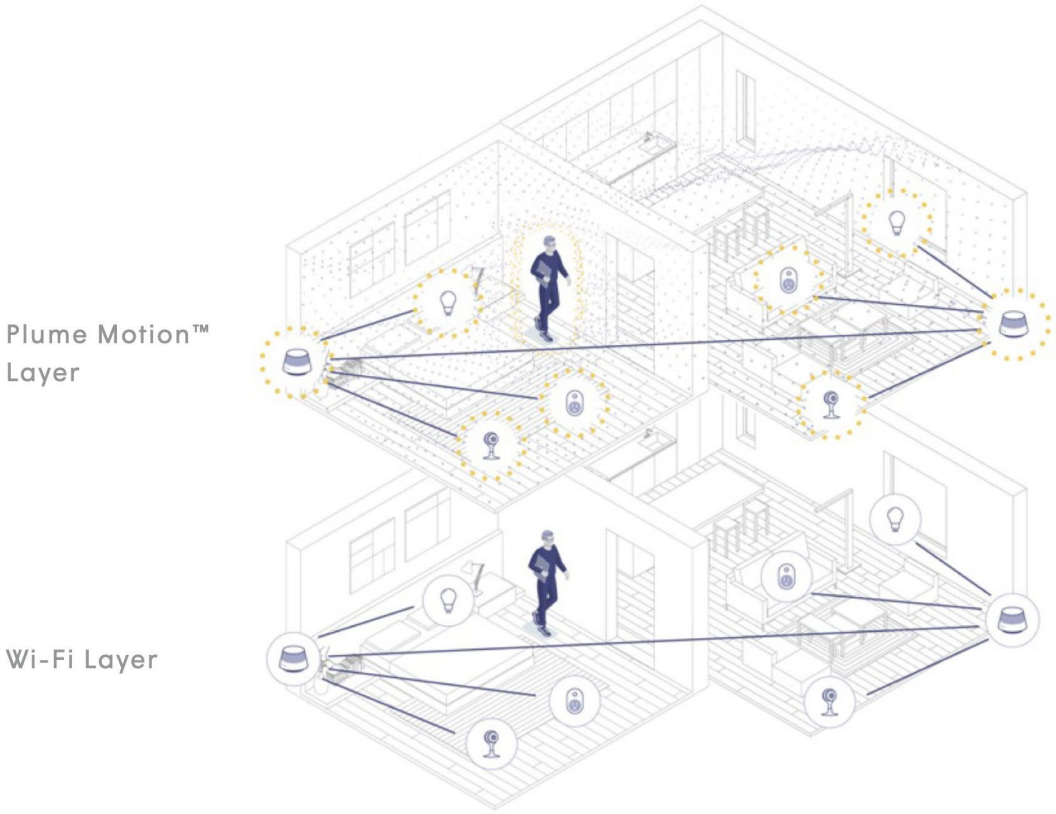

Figure 19 - Wi-Fi Motion Detection Concept

Using Wi-Fi to detect motion has a number of advantages:

- Far lower privacy concerns than using cameras in the home
- Does not require installation of extra sensors
- Does not require wearing of tags or sensors
- Comes with no hardware cost - can be implemented as a pure software solution on existing Wi-Fi gear

Wi-Fi motion can be used to provide a number of services of interest to the user. Home physical security is an immediate application. When out of the house, the system can be put into a mode where an alarm is sent if motion is sensed in the house. Elder care is another application, based on the reverse concept. In elder care, an alarm is sent if there is *not* motion across the appropriate periods of the day. Home energy management, for example turning on the heat when someone arrives home, is another application. Future versions of the technology may even be able to detect falls, and react quickly enough to control lights as you enter a room.

5. Conclusion

COVID-19 has changed the environment for service providers. It has dramatically increased the usage, loads, and sensitive traffic in homes, increasing the difficulty of providing sufficient quality to consumers. And, these changes are likely to persist even after COVID-19 has receded. However, technologies exist that have proven to mitigate many of the issues detailed in this paper. Multiple-AP deployments, with intelligent optimization and throughput optimized steering can ensure sufficient throughput to all locations in the home. Interference in the most challenging environments, multi-dwelling units, can be minimized through cloud based joint optimization. Increased cyber security attacks can be thwarted with machine learning based anomaly detection. And, accurate evaluation of networks requiring attention can be achieved using quality of experience metrics. COVID-19 is also affecting service providers revenue and profit. Artificial intelligence can be applied to perform proactive support, reducing the costs of service calls, truck rolls, and most critically, customer churn. And new technologies such as Wi-Fi motion detection can be used to create new services and new revenue streams.

Abbreviations

AP	access point
COVID-19	Coronavirus Disease 2019
DHCP	dynamic host configuration protocol
DNS	domain name system
Mb/s	Megabits per second
MDU	multi-dwelling unit
QoE	quality of experience
QoS	quality of service
UPNP	universal plug and play
WFH	working from home